We claim:

1. An edge node adapted to serve a multiplicity of data streams, said edge node comprising:

a plurality of input ports adapted to receive said data streams, wherein each of said input ports includes an input-port controller and at least one of said input ports includes a bitrate-estimation device adapted to compute a bitrate requirement for each of said data streams;

a plurality of output ports, wherein each of said output ports includes an output-port controller;

a switching fabric adapted to connect any of said input ports to any of said output ports; and

an edge controller including:

an edge control processor adapted to communicate with said plurality of input ports and said plurality of output ports;

a route selection device, in communication with said edge control processor, adapted to select a route for each of said data streams;

a fabric scheduling device, in communication with said edge control processor, adapted to determine a distinct time of transfer for each of a set of data segments, into which each of said data streams are segmented, across said switching fabric; and

a bitrate-allocation device, in communication with said edge control processor, adapted to allocate a bitrate of a data stream based, at least in part, on said bitrate requirement for said data stream computed by said bitrate-estimation device.

2. The edge node of claim 1 wherein said switching fabric has a greater number of output ports than input ports.

3. The edge node of claim 1 wherein at least one of said output ports includes a time locking device adapted to control a data transmit time from said at least one of said output ports.

4. The edge node of claim 1 wherein said switching fabric is a common-memory fabric.

5. The edge node of claim 1 wherein said switching fabric is a space-switching fabric.

6. The edge node of claim 1 wherein said switching fabric is a rotator-based fabric.

7. The edge node of claim 1 wherein said switching fabric switches entire channels.

8. The edge node of claim 1 wherein said switching fabric switches data segments of time-shared channels.

9. The edge node of claim 1 wherein at least one of said plurality of input ports is a source port adapted to receive data from data sources and at least one of said plurality of input ports is a receiving port adapted to receive data from core nodes and other edge nodes.

10. The edge node of claim 1 wherein at least one of said plurality of output ports is a sink port adapted to transmit data to data sinks and at least one of said plurality of output ports is a departure port adapted to transmit data to core nodes and other edge nodes.

11. The edge node of claim 10 wherein each of said plurality of output ports is adapted to collate said data segments.

12. The edge node of claim 11 wherein said departure port is a time-locking departure port, said time-locking departure port including data buffers and time-locking circuitry.

13. The edge node of claim 1 wherein said bitrate-estimation device is adapted to base said bitrate requirement for each of said data streams on parameters associated with said each of said data streams.

14. The edge node of claim 1 where said bitrate-estimation device further comprises a hysteresis-control device adapted to base said bitrate requirement for each of said data streams on usage measurements associated with said each of said data streams.

15. The edge node of claim 14 wherein said input-port controller further comprises an input buffer and said hysteresis-control device comprises:

an on-off hysteresis-control unit adapted to:

determine an admission state of either accept or reject based on an occupancy of said input buffer; and

reject a request for a new connection for a data stream when said admission state is reject;

an incremental hysteresis-control unit adapted to request bitrate allocation increments.

16. A method of determining a bitrate requirement for a connection, said method comprising:

receiving a packet, where said packet includes an indication of a type for said connection;

determining, from said indication, a type for said connection;

if said connection is determined to be of a first type, extracting an explicit bitrate-allocation requirement from said packet;

if said connection is determined to be of a second type, computing said bitrate requirement from parameters characterizing said traffic associated with said connection; and

if said connection type is determined to be a of third type, deducing said bitrate requirement from data stream monitoring.

17. The method of claim 16 wherein said data stream monitoring comprises:

observing a buffer to obtain successive buffer-occupancy readings;

retaining at least two said buffer-occupancy reading;

comparing said at least two consecutive readings; and

determining a requirement for a bitrate allocation change based on said comparing.

18. The method of claim 17 further comprising modifying said deduced bitrate requirement based on said determining.

19. An edge node for high-speed traffic processing comprising:

a switching fabric;

a plurality of input ports in communication with said switching fabric, each of said plurality of input ports having an input-port controller;

a plurality of output ports in communication with said switching fabric, each of said plurality of output ports having an output-port controller; and

an edge node controller in communication with said switching fabric, each of said input-port controllers and each of said output-port controllers, said edge node controller including a high-speed scheduling device, said high-speed scheduling device operable to:

receive bitrate allocation information from at least one input-port controller;

maintain a state of a given input port associated with said at least one input-port controller;

maintain a state of each of said plurality of output ports;

assign time slots of a scheduling frame to communicate data segments from said given input port to one or more of said output ports, where the number of said time slots allocated to a given output port is based on said bitrate allocation information; and

transmit said scheduling frame to said corresponding input-port controller.

20. The edge node of claim 19 wherein at least one of said output port controllers further includes a time-locking device for time-locking a corresponding one of said plurality of output ports to a core node.

21. The edge node of claim 19 wherein at least one of said input port controllers comprises a destination identifier adapted to translate an Internet Protocol address.

22. The edge node of claim 19 wherein at least one of said input port controllers includes a data formatter adapted to segment a received data stream into equal sized data segments.

23. The edge node of claim 19 wherein at least one of said input port controllers includes a bitrate allocation requirement estimator adapted to control requests for changes in bitrate allocation.

24. The edge node of claim 19 wherein said input-port controller is adapted to receive a data stream that specifies a destination and an explicit bitrate allocation requirement and wherein said edge node controller is adapted to provide a route to said destination having said explicit bitrate allocation requirement.

25. The edge node of claim 19 wherein said input-port controller is adapted to receive a data stream that specifies a destination and estimate a bitrate allocation requirement for said data stream and wherein said edge node controller is adapted to provide a route to said destination having said estimated bitrate allocation requirement.

26. The edge node of claim 25 wherein said input-port controller is adapted to

estimate a requirement for bitrate allocation according to a hysteresis-control mechanism; and

transmit a request for a change in said estimated bitrate allocation requirement for said data stream.

27. A method of data scheduling in an edge node comprising:

receiving bitrate allocation information from an input-port controller;

maintaining a state of a given input port associated with said input-port controller;

maintaining a state of each of a plurality of output ports;

assigning time slots of a scheduling frame to communicate data segments from said given input port to one or more of said output ports, where the number of said time slots allocated to a given output port is based on said bitrate allocation information; and

transmitting said scheduling frame to said input-port controller.

28. A device for high-speed scheduling operable to:

receive a bitrate allocation from an input-port controller;

generate a bitrate demand matrix from said bitrate allocation;

where a schedule is defined by a plurality of time slots, assign said time slots to input-output pairs according to said bitrate demand matrix, an input state matrix and an output port state matrix; and

transmit said schedule to said input-port controller.

29. A method of scheduling a transfer of data segments through a switching fabric from a plurality of input ports of said switching fabric to a plurality of output ports of said switching fabric wherein a scheduling frame having a predetermined number of

time slots and an equal, predetermined duration is associated with each of said plurality of input ports, said method comprising:

dividing said duration of each said scheduling frame into a plurality of time windows;

dividing said plurality of input ports into input-port groups;

for each said scheduling frame, forming a plurality of non-intersecting scheduling domains, where each said scheduling domain associates one of said input-port groups with one of said plurality of time windows;

concurrently executing scheduling processes, one scheduling process for each of said non-intersecting scheduling domains, each said scheduling process determining a schedule that indicates, for a given scheduling domain, a timing for a transfer of data segments to selected ones of said plurality of output ports from said input ports associated with said time window in said given scheduling domain; and

at least where said executing does not use all said time windows of a given scheduling frame, repeating said forming and said executing for permutations of said input-port groups and time windows.

30. The method of claim 29 wherein, for a scheduling frame associated with a given input port, said executing scheduling process includes allocating at least one time slot to an output port to which said given input port has non-zero traffic.

31. The method of claim 30 wherein said executing said scheduling process for a given scheduling domain comprises sequentially executing said scheduling process for said input ports of said input-port group of said given scheduling domain.

32. The method of claim 29 further comprising receiving, from each of said plurality of input ports, a demand vector, where said demand vector indicates a required number of time slots per scheduling frame for each of said plurality of output ports and wherein said demand vector is structured as an aggregate-demand ring-list including a plurality of data records, each said data record including a field indicating

a preceding record identifier, a field indicating a succeeding record identifier, a field indicating an output port identifier, and a field indicating a non-zero number of required time slots per scheduling frame.

33. The method of claim 29 further comprising receiving, from each of said plurality of input ports, a demand vector, where said demand vector indicates a required number of time slots per scheduling frame for each of said plurality of output ports and wherein said demand vector is structured as a flat-demand ring-list including a plurality of data records, each said data record including a field indicating a preceding record identifier, a field indicating a succeeding record identifier, and a field indicating an output port identifier.

34. The method of claim 29 further comprising limiting an execution time period for each said scheduling process.

35. The method of claim 34 wherein said limited execution time period varies with each of said repeating of said concurrently executing a scheduling process.

36. The method of claim 29 wherein said switching fabric is a space switch.

37. The method of claim 29 wherein said switching fabric is a rotator-based switch having a plurality of middle memories.

38. The method of claim 37 wherein, for each said input port, a time slot is uniquely associated with one of said plurality of middle memories.

39. The method of claim 29 further comprising using said schedule, determined for said duration of said scheduling frame, for a configuration period, where the duration of said configuration period is an integer multiple of said duration of scheduling frame.

40. In an edge node comprising a plurality of input ports, a plurality of output ports and a switching fabric, a fabric-scheduling apparatus comprising:

a plurality of input-port scheduling devices, where each of said input-port scheduling devices is associated with an input-port group and each input-port

group is a sub-set of said plurality of input ports, each of said input-port scheduling devices including:

a receiver adapted to receive bitrate allocations from controllers of said plurality of input ports and translate each said bitrate allocation into translated bitrate allocations, where each said translated bitrate allocation indicates a required number of time-slots for a predefined scheduling frame duration;

a ring-list generator adapted to form a ring-list from said translated bitrate allocations received from said receiver;

a ring-list time-slot assignment circuit adapted to:

receive said ring-list from said ring-list generator;

perform an assignment of each time-slot in a predefined time-slotted scheduling frame to a transfer of a data segment from one of said input ports in said input-port group associated with said input-port scheduling devices to one of said plurality of output ports; and

modify said ring-list to reflect said assignment;

a ring-list memory adapted to maintain said modified ring-list; and

a plurality of output-state memory devices, each of said output-state memory devices associated with a predefined time window; and

an access-control device adapted to cyclically connect each of said input-port scheduling devices with each of said output-state memory devices.

41. The edge node of claim 40 wherein said access control device comprises a two-way rotating connector.

42. The edge node of claim 41 wherein said ring-list generator generates an aggregate-demand ring-list.

43. The edge node of claim 41 wherein said ring-list generator generates a flat-demand ring-list.